

# Secure Anonymous Broadcast

Mahnush Movahedi  
movahedi@cs.unm.edu

Jared Saia  
saia@cs.unm.edu

Mahdi Zamani  
zamani@cs.unm.edu

Department of Computer Science, University of New Mexico, Albuquerque, NM, USA 87131

## Abstract

In anonymous broadcast, one or more parties want to anonymously send messages to all parties. This problem is increasingly important as a black-box in many privacy-preserving applications such as anonymous communication, distributed auctions, and multi-party computation. In this paper, we design decentralized protocols for anonymous broadcast that require each party to send (and compute) a polylogarithmic number of bits (and operations) per anonymous bit delivered with  $O(\log n)$  rounds of communication. Our protocol is provably secure against traffic analysis, does not require any trusted party, and is completely load-balanced. The protocol tolerates up to  $n/6$  statically-scheduled Byzantine parties that are controlled by a computationally unbounded adversary. Our main strategy for achieving scalability is to perform local communications (and computations) among a logarithmic number of parties. We provide simulation results to show that our protocol improves significantly over previous work. We finally show that using a common cryptographic tool in our protocol one can achieve practical results for anonymous broadcast.

## 1 Introduction

Today, political and commercial entities are increasingly engaging in sophisticated cyber-warfare to damage, disrupt, or censor information content [32]. In designing anonymous communication services, there is a need to ensure reliability even against very powerful types of adversaries. Such adversaries can monitor large portions of networks and control a certain fraction of the parties to run sophisticated *active* attacks such as jamming, corruption, and forging, as well as simple *passive* attacks such as eavesdropping and non-participation.

In this paper, we focus on the problem of secure anonymous broadcast, where a set of  $n$  parties want to anonymously send their messages to all parties. Anonymous broadcast is an important tool for achieving privacy in several distributed applications such as anonymous communication [10], private information retrieval [12], secure auctions [20], and multi-party computation (MPC). Our goal is to design a decentralized anonymous broadcast protocol that scales well with the number of parties and is robust against an active adversary. One motivating application for this protocol is a decentralized version of Twitter that enables provably-anonymous broadcast of messages.

One challenging problem with most anonymity-based systems is resistance against *traffic-analysis*. A global adversary can monitor traffic exchanged between parties to link messages to the their corresponding senders. Such a powerful adversary was assumed to be unrealistic in the past but it is believed to be realistic today especially if the service provider is controlled or compromised by a state-level surveillance authority [19]. Unfortunately, well-known anonymous services such as Crowds [39] and Tor [17] are not secure against traffic analysis attacks.

Two widely-accepted architectures for providing general anonymity against an active adversary are *Mix networks (Mix-Nets)* and *Dining Cryptographers networks (DC-Nets)*, both of which were

originally proposed by Chaum [10, 11]. Mix-Nets require semi-trusted infrastructure nodes and are known to be vulnerable to traffic analysis and active attacks [38]. DC-Nets [11, 23, 41, 42], on the other hand, provide anonymous broadcast protocols among a group of parties without requiring trusted parties. The core idea of DC-Nets is that a protocol for multi-party computation can be used to perform sender and receiver anonymous broadcast. For example, if party  $p_i$  wants to broadcast a message  $m_i$  anonymously, then all other parties participate in a multi-party sum with input zero, while party  $p_i$  participates with input  $m_i$ . All parties learn the sum, which is  $m_i$  while all inputs remain private. This ensures that no party can trace the output message  $m_i$  to its input, keeping  $p_i$  anonymous.

Although DC-Nets are provably-secure against traffic analysis, they face several challenges. First, a reservation mechanism is required to schedule which party is broadcasting without compromising the anonymity of the sender. Second, DC-Nets are susceptible to collisions, which degrade throughput. A jamming adversary may even use collisions to render the channel useless by continuously transmitting in every round. Third, typical DC-Nets are not scalable given that the bit complexity required to anonymously broadcast a single bit among  $n$  parties is  $\Omega(n^2)$ .

State-of-the-art approaches that address some of these challenges include [14, 23, 41]. The majority of these methods scale poorly with network size, rendering them impractical for large networks. Recently, Zamani et al. [44, 28] proposed the first anonymous broadcast protocol where each party sends  $o(n)$  bits to broadcast a bit among  $n$  parties. Their protocol uses multi-party computation to achieve full anonymity and logarithmic-size groups of parties to achieve  $\tilde{O}(1)$  communication and computation costs. Unfortunately, their protocol has polylogarithmic rounds of communication and is not practical due to large logarithmic factors hidden in the complexity notation.

To the best of our knowledge, every sender and receiver anonymous broadcast protocol that does not rely on a trusted party consists of at least three steps.

1. **Input:** Initially, each party holds a message. The party distributes its message or a representation of it among all or a subset of parties. This step requires sending  $\Omega(n)$  messages.
2. **Multi-party shuffling:** All or a subset of parties participate in a multi-party protocol to obliviously generate a random permutation of the sequence of message they hold.
3. **Output:** All or a subset of parties holding a sequence of messages broadcast them to all parties. This step requires sending  $\Omega(n^2)$  messages for delivering  $n$  shuffled messages.

Since much cannot be done to improve the cost of the output phase, we will mainly focus on the multi-party shuffling step in this paper. Multi-party shuffling can be used as a black-box in multi-party computation problems. Boyle et al. [7] use oblivious shuffling to randomly choose inputs for a sublinear function evaluation, where a function is evaluated over  $o(n)$  inputs chosen uniformly at random in the presence of an active adversary. Laur et al. [33] and Goodrich et al. [25] describe how multi-party shuffling can be used for implementing oblivious database operations and oblivious storage.

**Our Model.** In this paper, we consider a network of  $n$  parties whose identities are common knowledge. We assume there is a private and authenticated communication channel between every pair of parties and the communication is *synchronous*. Our protocol does not require the presence of any trusted third party, and we do not assume the existence of a reliable broadcast channel. We assume  $t < (1/6 - \epsilon)n$  of the parties are controlled by an *active* adversary, for some positive constant  $\epsilon$ . We assume our adversary is *computationally unbounded* and is actively trying to prevent the protocol from succeeding by attacking the privacy of the parties, and the integrity of communications, by attempting to corrupt, forge, or drop messages. We say that the parties controlled by the adversary are *dishonest* and that the remaining parties are *honest* meaning that they strictly follow our

protocol. We finally assume that the adversary is *static* meaning that it must select the set of dishonest parties at the start of the protocol.

### 1.1 Our Contribution

In this paper, we design a decentralized protocol for anonymous communication that is secure against an active adversary. Our protocol is efficient and scales well with the number of parties. Moreover, our protocol is load-balanced meaning that each party handles a roughly equal amount of communication and computation. We use techniques from multi-party computation, where a set of  $n$  parties, each having a secret value, compute a known function over their inputs, without revealing the inputs to any party.

Recently, Boyle et al. [7] and Dani et al. [16] proposed scalable solutions to general MPC. Unfortunately, both of these protocols are not practical due to large logarithmic and constant factors in their communication/computation costs. Moreover, the protocol of Boyle et al. is not load-balanced making it hard to be used in settings like mobile networks where the parties have limited resources. Despite their inefficiency, we are inspired by [7] and [16] to achieve scalability by performing local communications and computations in logarithmic-size groups of parties called *quorums*, where the number of dishonest parties in each quorum is guaranteed not to exceed a certain fraction. Using quorums and by simplifying much of their work, we develop an efficient multi-party shuffling protocol for anonymizing user inputs.

Our protocol is provably-secure as it is based on a formal security framework, which follows from the security of MPC. We show that the anonymity achieved by this method is, in particular, resistant to traffic analysis. We also provide provable anonymity against *a priori* knowledge that an adversary might have regarding the potential communicating parties. Moreover, unlike the majority of previous work which rely on centralized trusted servers, our protocol is fully-decentralized and does not require any trusted party.

Our protocol has polylogarithmic communication and computation costs with respect to the number of parties. We prove the following main theorem in Section A.

**Theorem 1.** *Consider  $n$  parties in a fully connected network with private channels where each party has a message to send to all parties. There exists an unconditionally-secure  $n$ -party protocol such that if all honest parties follow the protocol, then with high probability<sup>1</sup>:*

- *Each honest party sends its message to all parties anonymously.*
- *The protocol tolerates up to  $t < (1/6 - \epsilon)n$  malicious parties, for some positive constant  $\epsilon$ .*
- *Each party sends  $\tilde{O}(1)$  bits and computes  $\tilde{O}(1)$  operations for shuffling  $n$  messages<sup>2</sup>.*
- *Each party sends  $\tilde{O}(1)$  bits and computes  $\tilde{O}(1)$  operations for delivering one anonymous bit.*
- *The protocol has  $O(\log n)$  rounds of communication.*

**Protocol Overview.** In our protocol,  $n$  parties participate in a MPC to jointly compute a shuffling function that randomly permutes their messages. Then, the results are broadcast to all participating parties. More specifically, our protocol builds a set of quorums in a one-time setup phase and then, uses the quorums in the online phase for shuffling input messages. We represent the desired shuffling function by an arithmetic circuit, where the computation of each gate is assigned to a quorum. Then, we evaluate the circuit level-by-level, passing the outputs of one level as the inputs

<sup>1</sup>An event occurs *with high probability*, if it occurs with probability  $\geq 1 - 1/n^c$ , for  $c > 0$  and sufficiently large  $n$ .

<sup>2</sup>The  $\tilde{O}$  notation is used as a variant of the big-O notation that ignores logarithmic factors. Thus,  $f(n) = \tilde{O}(g(n))$  means  $f(n) = O(g(n) \log^k g(n))$  for some  $k$ .

to the next level. Once the local computation is finished in each quorum, the result is forwarded to the next quorum via one-to-one communication with parties of the next quorum. Finally, at the highest level, the shuffled messages are computed and sent to all parties.

One technique for randomly shuffling a set of messages in a multi-party setting is to assign to each message a uniform random value and then, obviously sort the messages according to the random values. One issue with this technique is that if the randoms are not distinct (i.e., there is a *collision*), then the resulting distribution deviates from the uniform distribution. Unfortunately, this gives the adversary some advantage to map the inputs to outputs and thus, break the anonymity. More formally, let  $n$  be the number of messages and  $[1, M]$  be the range of random values we choose for each message. With  $n$  such random values, we can generate  $M^n$  possible states while there are  $n!$  possible permutations of the messages. Since  $n! \nmid M^n$ , it is easy to see that some permutations are more likely to be generated than others. On the other hand, we will show that by choosing a sufficiently large  $M$ , we can prevent collisions with high probability. Using the Chernoff bound, we will prove that choosing  $M = \Omega(kn^2 \log n)$  guarantees a uniform random shuffle with probability  $1 - 1/n^k$ , for any constant  $k > 0$ .

## 2 Further Related Work

Some protocols are built upon a relaxed notion of anonymity called *k-anonymity* [41, 43, 34], where the adversary is assumed to be unable to identify the actual sender/receiver of a message from a set of  $k$  parties (called *anonymity set*). Even though *k-anonymity* often increases efficiency significantly, choosing small  $k$ 's can result in severe privacy problems. For example, attackers often have background knowledge and it is shown that small anonymity sets are likely to leak privacy when attackers have such knowledge [36]. For example, a person located in New Mexico is more likely to search for a restaurant serving chili stew than a person in Vermont.

Von Ahn et al. [41] develop a cryptographic broadcast protocol based on DC-Nets that is resistant to a static active adversary. A set of  $n$  parties with private inputs compute and share the sum of their inputs without revealing any parties' input. The authors introduce *k-anonymity*, which means no polynomial-time adversary may distinguish the sender/receiver of a message from among  $k$  honest senders/receivers. To achieve *k-anonymity*, they partition the set of parties into groups of size  $M = O(k)$  and execute a multi-party sum protocol inside each group. The jamming detection mechanism is weak against an adversary who may waste valuable resources by adaptively filling up to  $M$  channels. In the case where *n-anonymity* is desired, the protocol requires  $O(n^3)$  messages to be sent per anonymous message and the total bit complexity is  $O(n^4)$ . The protocol has latency that is  $O(1)$  on average when the number of broadcasts is large, but which can be  $O(n)$  in worst case for a single broadcast.

Golle and Juels [23] employ cryptographic proofs of correctness to solve the jamming problem in DC-Nets assuming a static Byzantine adversary. The protocol detects jamming with high probability in  $O(1)$  rounds, requiring a total communication and computation complexity of  $O(n^2)$  bits. Their protocol assumes the existence of a reliable broadcast and a centralized trusted authority for key management distribution.

The Xor-trees approach of [18] extends DC-Nets to achieve  $O(n)$  amortized bit complexity, which is optimal. In this protocol, only a single user is allowed to send at any one time in a Xor-tree. Hence, the protocol is subject to performance degradation due to collisions as the number of users increases. The protocol assumes the existence of a public-key infrastructure and a non-Byzantine polynomial-time adversary. The total bit complexity of the protocol is  $O(n^2 t^2)$  bits in worst case, where  $t$  is the number of dishonest parties. The latency of the protocol is  $O(n)$  in worst case. However, a sender may broadcast large payloads to amortize the costs. The amortized

latency of the protocol is  $O(1)$ .

The Verdict protocol of [14] (which is based on Dissent [13]) has a client-server architecture and uses verifiable DC-Nets, where participants use public-key cryptography to construct ciphertext, and knowledge proofs to detect and exclude jamming parties before disruption. The protocol assumes the existence of a few highly-available servers, where at least one server is honest. All servers must be alive, however, for the protocol to work. An interesting aspect of Verdict is that it is robust to a large fraction of Byzantine parties (up to  $n - 2$ ). The paper demonstrates empirically that the system scales well with the number of parties, when the number of servers is fixed. The Tarzan protocol of Freedman and Morris [21] provides resistance against traffic analysis, but only against a passive adversary.

The Aqua protocol of Le Blond et al. [34] provides  $k$ -anonymity with traffic-analysis resistance against passive global attacks and active local attacks. The protocol achieves anonymity in a way similar to Tor (onion routing) and achieves unobservability through traffic obfuscating, which is to add artificial delay or artificial traffic (called chaff) to the connection.

In the last three decades, a large body of work has been devoted to designing MPC protocols [5, 4, 22, 15]. Unfortunately, most of these protocols are inefficient and scale poorly with the number of parties. Several MPC sorting and shuffling protocols have been proposed in the literature [24, 33, 45, 26]. Laur et al. [33] describe a multi-party shuffling protocol that can be used for anonymizing a set of inputs. Although the communication and round complexity of their protocol scales well with the number of inputs, they scale exponentially with the number of parties and hence, the method cannot be used in our model, where  $n$  is relatively large.

Goodrich [24] proposes an efficient data-oblivious randomized shellsort algorithm. Unfortunately, when implemented in a multi-party setting, this protocol requires  $O(m)$  rounds of communication to sort  $m$  values and has communication complexity  $O(\ell^2 n^2 m \log m)$ , where  $\ell$  is the message size. Zhang [45] and Hamada et al. [26] develop constant-round MPC sorting protocols that scale well with the number of inputs but scale poorly with the number of parties.

### 3 Preliminaries

In this section, we define standard terms, notation, and results used throughout the paper.

#### 3.1 Notation

An event occurs *with high probability*, if it occurs with probability at least  $1 - 1/n^c$ , for any  $c > 0$  and all sufficiently large  $n$ . We denote the set of integers  $\{1, \dots, n\}$  by  $[n]$ . Also, let  $\mathbb{Z}_p$  denote the additive group of integers modulo a prime  $p$ .

#### 3.2 Basic Tools

In this section, we review the definitions of standard basic tools used throughout the paper.

**Verifiable Secret Sharing.** An  $(n, t)$ -*secret sharing* scheme, is a protocol in which a dealer who holds a secret value shares the secret among  $n$  parties such that any set of  $t$  parties cannot gain any information about the secret, but any set of at least  $t + 1$  parties can reconstruct it. An  $(n, t)$ -*verifiable secret sharing (VSS)* scheme is an  $(n, t)$ -secret sharing scheme with the additional property that after the sharing phase, a dishonest dealer is either disqualified or the honest parties can reconstruct the secret, even if shares sent by dishonest parties are spurious. In our protocol, we use the constant-round VSS protocol of Katz et al. [30] that is based on Shamir’s secret sharing scheme [40]. This result is described in Theorem 2.



**Theorem 2.** [30] *There exists a synchronous linear  $(n, t)$ -VSS scheme for  $t < n/3$  that is perfectly-secure against a static active adversary. The protocol requires one broadcast and three rounds of communication.*

For practical purposes, one can use the cryptographic VSS scheme of Kate et al. [29] called *eVSS* (stands for efficient VSS), which is based on Shamir’s scheme and the hardness of the Discrete Logarithm (DL) problem. Since eVSS generates commitments over elliptic curve groups, it requires smaller message sizes than other DL-based VSS scheme such as [22].

**Theorem 3.** [29] *There exists a synchronous linear  $(n, t)$ -VSS scheme for  $t < n/2$  that is secure against a computationally-bounded static adversary. In worst case, the protocol requires two broadcasts and four rounds of communication.*

**Quorum Building.** A *good quorum* is a set of  $N = O(\log n)$  parties that contains a majority of honest parties. King et al. [31] showed that a *Byzantine Agreement (BA)* protocol can be used to bring all parties to agreement on a collection of  $n$  good quorums. In this paper, we use the fast BA protocol of Braud-Santoni et al. [8] to build  $n$  good quorums.

**Theorem 4.** [31, 8] *There exists a constant-round unconditionally-secure protocol that brings all good parties to agreement on  $n$  good quorums with high probability. The protocol has  $\tilde{O}(n)$  amortized communication and computation complexity<sup>1</sup>, and it can tolerate up to  $t < (1/3 - \epsilon)n$  malicious parties.*

**Secure Broadcast.** In the malicious setting, when parties have only access to secure pairwise channels, a protocol is required to ensure secure (reliable) broadcast<sup>2</sup>. Such a broadcast protocol guarantees all parties receive the same message even if the broadcaster (dealer) is dishonest and sends different messages to different parties. It is known that a BA protocol can be used to perform secure broadcasts. In our protocol, we use the BA algorithm of Braud-Santoni et al. [8] to perform secure broadcasts.

**Theorem 5.** [8] *There exists a constant-round unconditionally-secure protocol for performing secure broadcasts among  $n$  parties. The protocol has  $\tilde{O}(n)$  amortized communication and computation complexity, and it can tolerate up to  $t < (1/3 - \epsilon)n$  malicious parties.*

**Sorting Networks.** A *sorting network* is a network of *comparators*. Each comparator is a gate with two input wires and two output wires. When two values enter a comparator, it outputs the lower value on the top output wire, and the higher value on the bottom output wire. Ajtai et al. [1] describe an asymptotically-optimal  $O(\log n)$  depth sorting network. However, this network is not practical due to large constants hidden in the depth complexity. Leighton and Plaxton [35] propose a practical *probabilistic sorting circuit* that sorts with very high probability meaning that it sorts all but  $\epsilon \cdot n!$  of the  $n!$  possible input permutations, where  $\epsilon = 1/2^{2^{k\sqrt{\log n}}}$ , for any constant  $k > 0$ . For example, for  $k = 2$  and  $n > 64$ , we get  $\epsilon < 10^{-9}$ . While this circuit is sufficient for us to prove our main results (Theorem 1), one can instead use the  $O(\log^2 n)$ -depth sorting network of Batcher [3] for sorting all  $n!$  permutations at the expense of  $O(\log^2 n)$  protocol latency.

**Secure Comparison.** Given two linearly secret-shared values  $a, b \in \mathbb{Z}_p$ , Nishide and Ohta [37] propose an efficient protocol for computing a sharing of  $\rho \in \{0, 1\}$  such that  $\rho = (a \leq b)$ . Their protocol has  $O(1)$  rounds and requires  $O(\ell)$  invocations of a secure multiplication protocol, where

<sup>1</sup>Amortized communication complexity is the total number of bits exchanged divided by the number of parties.

<sup>2</sup>We are not aware of any easy approach to physically implement a broadcast channel without assuming a trusted party.

$\ell$  is the bit-length of elements to be compared. We refer to this protocol by **Compare**. We also describe a fast multiplication protocol to be used along with the comparison protocol of [37] for implementing fast comparator gates.

**Share Renewal.** In our protocol, a shuffling circuit is securely evaluated. Each gate of the circuit is assigned a quorum  $Q$  and the parties in  $Q$  are responsible for comparison of secret-shared inputs. Then, they send the secret-shared result to any quorums associated with gates that need this result as input. Let  $Q'$  be one such quorum. In order to secret-share the result to  $Q'$  without revealing any information to any individual party (or to any coalition of dishonest parties), a fresh sharing of the result must be distributed in  $Q'$ . To this end, we use the share renewal protocol of Herzberg et al. [27]<sup>1</sup>. To update a shared value  $s$  defined over a degree  $d$  polynomial  $\phi(x)$ , the protocol generates a degree  $d$  random polynomial  $\delta(x)$  such that  $\delta(0) = 0$ . The new polynomial  $\phi^*(x)$  is then computed from  $\phi^*(x) = \phi(x) + \delta(x)$ . Since  $\phi^*(0) = \phi(0) + \delta(0) = s$ , the new polynomial  $\phi^*(x)$  defines a fresh sharing of  $s$ . Combined with the VSS scheme of Theorem 2 (or Theorem 3), this protocol is secure against an active adversary (with  $t < n/3$ ), takes constant number of rounds, and requires each party to send  $O(n^2)$  field elements. We refer to this protocol by **RenewShares** throughout this paper.

## 4 Our Protocol

In this section, we describe our protocol for anonymous broadcast. Consider  $n$  parties  $P_1, P_2, \dots, P_n$  each having a message  $x_i \in \mathbb{Z}_p$ , for a prime  $p$  and all  $i \in [n]$ . The parties want to anonymously send the messages to each other and receive the results back. We first describe an ideal functionality of our protocol where a hypothetical trusted party  $P$  computes the desired protocol outcome by communicating with all parties<sup>2</sup>. In every run of the protocol,  $P$  executes a shuffling protocol over the messages and sends the shuffled sequence of messages to all parties. The shuffling protocol first chooses a uniform random number  $r_i \in \mathbb{Z}_p$  to form an input pair  $(r_i, x_i)$  for each party  $P_i$  and for all  $i \in [n]$ . The protocol then uses a shuffling circuit that is based on the sorting network of Leighton and Plaxton [35] to sort the set of pairs  $\{(r_1, x_1), \dots, (r_n, x_n)\}$  according to their first elements. We later show that this functionality randomly permutes the set of inputs  $\{x_1, \dots, x_n\}$ .

In our protocol, we denote the shuffling circuit by  $C$ , which has  $m$  gates. Each gate is essentially a comparator gate and is denoted by  $G_i$ , for  $i \in [m]$ . Our protocol first creates  $n$  quorums and then assigns each gate of  $C$  to a quorum. For  $n$  parties, the circuit has  $\lceil n/2 \rceil$  input gates as each comparator gate has two inputs. We label the quorums assigned to the input gates by  $Q_1, \dots, Q_{\lceil n/2 \rceil}$  and call them *input quorums*.  $C$  also has  $\lceil n/2 \rceil$  output gates, which correspond to *output quorums* labeled by  $Q'_1, \dots, Q'_{\lceil n/2 \rceil}$ .

We now implement the real functionality of our protocol based on the ideal functionality described above. Protocol 1 defines our main algorithm. Throughout the protocol, we represent each shared value  $s \in \mathbb{Z}_p$  by  $\langle s \rangle = (s_1, \dots, s_n)$  meaning that each party  $P_i$  holds a share  $s_i$  generated by the VSS scheme during its sharing phase. Using the natural component-wise addition of representations, we define  $\langle a \rangle + \langle b \rangle = \langle a + b \rangle$ . For multiplication, we define  $\langle a \rangle \cdot \langle b \rangle = \text{Multiply}(\langle a \rangle, \langle b \rangle)$ , where **Multiply** is a protocol defined in this section.

For the reader's convenience, we now list all subprotocols that we use in our protocol. These subprotocols are also defined in this section, Section 3, and Section B.

- **GenRand:** A well-known technique (due to [4]) that generates uniform random secrets by adding shared values chosen uniformly at random by each party. Defined as Protocol 3 in Section B.

<sup>1</sup>This technique was first used by Ben-Or et al. [5] and was later proved to be UC-secure by Asharov and Lindell in [2].

<sup>2</sup>We are inspired by the standard ideal/real world definition for multi-party protocols proposed by Canetti [9].

- **Compare:** The protocol of Nishide and Ohta [37] for securely comparing two secret-shared values.
- **Multiply:** Multiplies two secret-shared values. Defined as Protocol 2 in this section.
- **RenewShares:** The share renewal protocol of Herzberg et al. [27] for re-randomizing shares of a secret-shared value. This makes past knowledge of the adversary obsolete after re-sharing the secret.
- **Reconst:** Given a shared secret  $s$ , reconstructs  $s$  via polynomial interpolation and error correction. Defined as Protocol 4 in Section B.

---

**Protocol 1 AnonymousBroadcast**


---

1. **Setup:** Parties jointly run the quorum building protocol of Theorem 4 with all parties to agree on  $n$  good quorums  $Q_1, \dots, Q_n$ , and assign each gate  $G_i$  to the  $(i \bmod n)$ -th quorum, for all  $i \in [m]$ .
  2. **Input Sharing:** Parties  $P_i$  and  $P_{i+1}$  secret share their inputs  $x_i$  and  $x_{i+1}$  among all parties of  $Q_{[i/2]}$  using the VSS scheme of Theorem 2, for all  $i \in [n]$ .
  3. **Random Generation:** Parties in each input quorum  $Q_{[i/2]}$  run GenRand twice to generate sharings  $\langle r_i \rangle$  and  $\langle r_{i+1} \rangle$  of two random elements  $r_i, r_{i+1} \in \mathbb{Z}_p$ , for some prime  $q$ . At the end of this step,  $Q_{[i/2]}$  holds two pairs of sharings  $(\langle r_i \rangle, \langle x_i \rangle)$  and  $(\langle r_{i+1} \rangle, \langle x_{i+1} \rangle)$ , for all  $i \in [n]$ .
  4. **Circuit Computation:** The circuit is evaluated level-by-level starting from the input gates. For each gate  $G$  and the quorum  $Q$  associated with it, parties in  $Q$  do the following.
    - (a) **Gate Computation:** Let  $(\langle r \rangle, \langle x \rangle)$  and  $(\langle r' \rangle, \langle x' \rangle)$  be the sharings associated with the inputs of  $G$ . Parties compute  $\langle \rho \rangle = \text{Compare}(\langle r \rangle, \langle r' \rangle)$ , where  $\rho = (r \leq r')$ . Then, they compute the output pairs  $(\langle s \rangle, \langle y \rangle)$  and  $(\langle s' \rangle, \langle y' \rangle)$  from
 
$$\begin{aligned} \langle s \rangle &= \langle \rho \rangle \cdot \langle r \rangle + (1 - \langle \rho \rangle) \cdot \langle r' \rangle & \langle s' \rangle &= \langle \rho \rangle \cdot \langle r' \rangle + (1 - \langle \rho \rangle) \cdot \langle r \rangle \\ \langle y \rangle &= \langle \rho \rangle \cdot \langle x \rangle + (1 - \langle \rho \rangle) \cdot \langle x' \rangle & \langle y' \rangle &= \langle \rho \rangle \cdot \langle x' \rangle + (1 - \langle \rho \rangle) \cdot \langle x \rangle \end{aligned} \quad (1)$$
 For every addition over two shared values  $\langle a \rangle$  and  $\langle b \rangle$  performed above, parties compute  $\langle c \rangle = \langle a \rangle + \langle b \rangle$ . For every multiplication, they run  $\langle c \rangle = \text{Multiply}(\langle a \rangle, \langle b \rangle)$ .
    - (b) **Output Resharing:** Parties run RenewShares over  $\langle s \rangle, \langle y \rangle, \langle s' \rangle$ , and  $\langle y' \rangle$  to re-share them in the quorum associated with the parent gate.
  5. **Output Propagation:** Let  $(\langle s_i \rangle, \langle y_i \rangle)$  and  $(\langle s_{i+1} \rangle, \langle y_{i+1} \rangle)$  be the pairs of shared values each output quorum  $Q'_i$  receives. Parties in  $Q'_i$  run  $z_i = \text{Reconst}(\langle y_i \rangle)$  and  $z_{i+1} = \text{Reconst}(\langle y_{i+1} \rangle)$  and send  $(z_i, z_{i+1})$  to all  $n$  parties. Each party receiving a set of  $N$  pairs from each output quorum, chooses one pair via majority filtering and considers it as the output of that quorum.
- 

We prove the correctness and secrecy of Protocol 1 (and Theorem 1) in Section A.2. In Lemma 4, we show that for sufficiently large  $k > 0$  and  $p > \frac{3}{2}kn^2 \log n$ , this protocol computes a random permutation of the input messages with high probability. We also prove the following lemmas in Section A.2.1.

**Lemma 1.** Using the perfectly-secure VSS of Theorem 2, Protocol 1 sends  $\tilde{O}(n)$  bits, computes  $\tilde{O}(n)$  operations, and takes  $O(\log n)$  rounds of communication for shuffling  $n$  messages.

**Lemma 2.** Using the cryptographic VSS of Theorem 3, Protocol 1 sends  $O(n \log^5 n)$  messages of size  $O(\kappa + \log n)$  and computes  $O(n(\kappa + \log n) \log^5 n)$  operations for shuffling  $n$  messages, where  $\kappa$  is the security parameter. The protocol has  $O(\log n)$  rounds of communication.



We now describe the subprotocol **Multiply** that is based on a well-known technique proposed by Beaver [4]. The technique generates a shared multiplication triple  $(\langle u \rangle, \langle v \rangle, \langle w \rangle)$  such that  $w = u \cdot v$ . The triple is then used to convert multiplications over shared values into additions. Since the triple generation step is independent of the inputs, it can be done in a preprocessing phase. Moreover, since each triple acts like a one-time pad, a fresh triple must be used for each multiplication.

---

**Protocol 2 Multiply**

---

*Usage.* Initially, parties hold two secret-shared values  $\langle a \rangle = (a_1, \dots, a_N)$  and  $\langle b \rangle = (b_1, \dots, b_N)$  that are on a polynomial of degree  $N/3$ . The protocol computes a shared value  $\langle c \rangle = (c_1, \dots, c_N)$  such that  $c = a \cdot b$ .

**Multiply** $(\langle a \rangle, \langle b \rangle)$  :

For all  $i \in [N]$ , party  $P_i$  performs the following steps.

1.  $P_i$  runs **GenRand** twice with other parties to generate two shared random values  $u_i$  and  $v_i$  both on polynomials of degree  $N/6$ . Then, he computes  $w_i = u_i \cdot v_i$
  2.  $P_i$  updates  $u_i$ ,  $v_i$ , and  $w_i$  by running **RenewShares**.
  3.  $P_i$  computes  $\varepsilon_i = a_i + u_i$  and  $\delta_i = b_i + v_i$  and runs  $\varepsilon = \text{Reconst}(\varepsilon_i)$  and  $\delta = \text{Reconst}(\delta_i)$ . Then, he computes  $c_i = w_i + \delta a_i + \varepsilon b_i - \varepsilon \delta$ .
- 

Clearly,  $\varepsilon$  and  $\delta$  can be safely revealed to all parties so that each party can compute  $\varepsilon \delta$  locally. The correctness and the secrecy of the algorithm are proved by Beaver [4]. The only difference between Beaver's original method and our protocol is as follows. Following [5], he proposes to generate shared random elements  $u$  and  $v$  on a polynomial of degree  $N/3$  and multiply them to get a polynomial of degree  $2N/3$  for  $w$ . Then, a degree reduction algorithm is run to reduce the degree from  $2N/3$  to  $N/3$ . Instead of this, we choose polynomials of degree  $N/6$  for  $u$  and  $v$  to get a polynomial of degree  $N/3$  for  $w$ . In our protocol, since up to  $N/6$  of parties are dishonest in each quorum, we can do this without revealing any information to the adversary. We prove this formally in Lemma 5 and Lemma 6.

One issue with multiplying shares, as noticed by Ben-Or et al. [5], is that the resulting polynomial is not completely random because, for example, the product of two polynomials cannot be irreducible. To solve this, we simply re-randomize the shares via **RenewShare**. Moreover, since  $\langle u \rangle$  and  $\langle v \rangle$  are on degree  $N/6$  polynomials and **Multiply** uses them to mask-and-reveal the inputs (that are on degree  $N/3$  polynomials) in the last step, it is necessary to randomize  $\langle u \rangle$  and  $\langle v \rangle$  via **RenewShares** before masking. This puts the shared values  $\langle u \rangle$  and  $\langle v \rangle$  on new random polynomials with degree  $N/3$ .

**Definition 1. [Perfect Random Permutation]** Consider a set of  $n \geq 1$  elements  $A = \{a_1, a_2, \dots, a_n\}$ . A perfect random permutation of  $A$  is a permutation chosen uniformly at random from the set of all possible  $n!$  permutations of  $A$ .

**Lemma 3.** *Consider a sequence of input pairs  $(r_1, x_1), \dots, (r_n, x_n)$ , and a sorting protocol  $\Pi$  that sorts the pairs according to their first elements.  $\Pi$  computes a perfect random permutation of the pairs if their first elements are chosen uniformly at random and are distinct.*

*Proof.* Let  $X = (r_1, x_1), \dots, (r_n, x_n)$  be the input sequence and  $Y = (s_1, y_1), \dots, (s_n, y_n)$  be the output sequence of  $\Pi$ . Note  $s_1, \dots, s_n$  is the sorted sequence of  $\{r_1, \dots, r_n\}$ . An arbitrary output sequence of pairs  $Y' = (s'_1, y'_1), \dots, (s'_n, y'_n)$  is said to be equal to  $Y$  if  $y_i = y'_i$ , for all  $i \in [n]$ . We want to prove that the probability of  $Y'$  being equal to  $Y$  is  $\frac{1}{n!}$ . In general, for any  $i \in [n]$ ,  $y_i = y'_i$  if and only if  $s'_i$

is the  $i$ -th smallest element in  $\{r_1, \dots, r_n\}$  conditioned on knowing the  $i-1$  smallest elements, which happens with probability  $\frac{1}{n-i+1}$ . Thus, the probability that  $Y = Y'$  is  $\frac{1}{n} \cdot \frac{1}{n-1} \cdot \dots \cdot \frac{1}{2} \cdot 1 = \frac{1}{n!}$ .  $\square$

In the random generation step of Protocol 1, it is possible that two or more input quorums choose the same random elements from  $\mathbb{Z}_p$ . In this situation, we say a *collision* happens. Collisions reduce the level of anonymity our protocol guarantees because the higher the probability of collisions, the higher the chance the adversary is given in guessing the correct sequence of inputs. On the other hand, we observe that if the field size (i.e.,  $p$ ) is sufficiently large, then the probability of collisions becomes overwhelmingly small. Lemma 4 gives a lower bound on  $p$  such that collisions are guaranteed to happen with negligible probability.

**Lemma 4.** *Let  $\mathbb{Z}_p$  be the field of random elements generated in the random generation step of Protocol 1. The probability there is a collision between any two parties is negligible if  $p > \frac{3}{2}kn^2 \log n$ , for some  $k > 0$ .*

*Proof.* Based on Theorem 4, all input quorums are good. Based on the correctness of GenRand, all elements generated by the input quorums in the random generation step of Protocol 1 are chosen uniformly at random and independent of all other random elements generated throughout the protocol. Let  $P_i$  and  $P_j$  be two parties and  $r_i$  and  $r_j$  be the random values assigned to them respectively by their corresponding input quorums. The probability that  $r_i = r_j$  is  $1/p$ . Let  $X_{ij}$  be the following indicator random variable and  $Y$  be a random variable giving the number of collisions between any two parties,

$$X_{ij} = \begin{cases} 1, & r_i = r_j \\ 0, & \text{otherwise} \end{cases}, \quad Y = \sum_{i,j \in [n]} X_{ij}.$$

Using the linearity of expectations,

$$E(Y) = E\left(\sum_{i,j \in [n]} X_{ij}\right) = \sum_{i,j \in [n]} E(X_{ij}) = \frac{1}{p} \binom{n}{2} = \frac{n(n-1)}{2p}.$$

We want to find an upper bound on the probability of collisions using the Chernoff bound defined by

$$\Pr(Y \geq (1 + \alpha)E(Y)) \leq e^{-\frac{\alpha^2 E(Y)}{3}}.$$

To ensure that no collision happens with high probability, we need to have  $(1 + \alpha)E(Y) < 1$  while  $e^{-\frac{\alpha^2 E(Y)}{3}} < \frac{1}{n^k}$ , for any  $k > 0$ . Choosing  $\alpha < \frac{1}{E(Y)} - 1$  and solving the inequalities for  $E(Y)$  we get

$$\begin{aligned} e^{-\frac{\alpha^2 E(Y)}{3}} < \frac{1}{n^k} &\Rightarrow e^{-\frac{\alpha^2 E(Y)}{3}} < e^{-k \log n} \Rightarrow 1 - \frac{\alpha^2 E(Y)}{3} < -k \log n \Rightarrow \frac{(\alpha+1)^2 E(Y)}{3} > k \log n \\ \frac{1}{3E(Y)} > k \log n &\Rightarrow E(Y) < \frac{1}{3k \log n} \end{aligned}$$

Since  $E(Y) = \frac{n(n-1)}{2p} < \frac{1}{3k \log n}$ , solving this for  $p$  gives the bound  $p > \frac{3}{2}kn^2 \log n$  and  $\alpha < 3k \log n - 1$ .  $\square$

## 5 Simulation Results

To study the feasibility of our scheme and compare it to previous work, we implemented an experimental simulation of our protocol and three other protocols which can be used for shuffling  $n$  inputs randomly (with traffic-analysis resistance) in the same setting. These protocols are due to Dani et al. [16], Boyle et al. [7], and Zamani et al. [44]. To the best of our knowledge, these protocols

have the best scalability with respect to the network size among other works described in the literature for the same setting. Since the protocols of [16] and [7] are general MPC algorithms, we use them for running our shuffling technique described in Section 4. The protocol of [44] is a jamming-resistant version of DC-Nets that scales better than other DC-Net protocols [11, 23, 41, 42]. We stress that we are interested in evaluating our protocols for large network sizes and hence, our choice of protocols for this section is based on their scalability for large  $n$ 's. One may find other protocols in the literature that perform better than the protocols of our choice for small  $n$ 's.

We run our protocol for inputs chosen from  $\mathbb{Z}_p$  with a 160-bit prime  $p$  for getting about 80 bits of security. We set the parameters of our protocol in such a way that we ensure the probability of error for the quorum building algorithm of [8] is smaller than  $10^{-5}$ . For the sorting circuit, we set  $k = 2$  to get  $\epsilon < 10^{-8}$  for all values of  $n$  in the experiment. Clearly, for larger values of  $n$ , the error becomes superpolynomially smaller, e.g., for  $n = 2^{25}$ , we get  $\epsilon < 10^{-300}$ . For all protocols evaluated in this section, we assume cheating (by malicious parties) happens in every round of the protocols. This is essential for evaluating various strategies used by these protocols for tolerating active attacks.

Figure 1 illustrates the simulation results obtained for various network sizes between  $2^5$  and  $2^{30}$  (i.e., between 32 and about 1 billion). To better compare the protocols, the vertical and horizontal axis of the plot are scaled logarithmically. The  $x$ -axis presents the number of parties and the  $y$ -axis presents the number of Kilobytes sent by each party for delivering one anonymous bit. In this figure, we report results from three different versions of our protocols. The first plot (marked with circles) belongs to our unconditionally-secure protocol that uses the perfectly-secure VSS scheme of Katz et al. [30]. The second plot (marked with stars) represents our computationally-secure protocol which uses the cryptographic VSS of Kate et al. [29]. The last plot (marked with diamonds) shows the cost of the cryptographic protocol with amortized (averaged) setup cost. To obtain the amortized plot, we run the setup phase of Protocol 1 once and then used the setup data to run our online protocol 100 times. The total number of bits sent was then divided by 100 to get the average communication cost. To achieve better results, we also generated a sufficient number of random triples in the setup phase. Then, the triples were used by our multiplication subprotocol in the online phase to multiply secret-shared values efficiently.

We observe that our protocols (even the unconditional version) perform significantly better than the other protocols. For example, for  $n = 2^{20}$  (about 1 million parties<sup>1</sup>), the amortized protocol requires each party to send about 64KB of data per anonymous bit delivered (about 8MB for our crypto version and about 64MB for our unconditional version) while the protocols of [7], [16], and [44] each send more than one Terabytes of data per party and per anonymous bit delivered.

## 6 Conclusion and Open Problems

We described an anonymous broadcast protocol that is fully decentralized and tolerates up to  $n/6$  active faults. Moreover, our protocol is load-balanced and can tolerate traffic-analysis attacks. The amount of information sent and the amount of computations performed by each party scales polylogarithmically with the number of parties. The scalability is achieved by performing local communications and computations in groups of logarithmic size and by relaxing the latency requirements.

Several open problems remain. First, can we decrease the number of rounds of our protocol using a smaller-depth sorting circuit? For example, since our protocol sorts uniform random numbers, it seems possible to use a smaller depth non-comparison-based sorting circuit like bucket sort.

---

<sup>1</sup>This is less than 1% of the number of active Twitter users. An intriguing application of our protocol is an anonymous version of Twitter.




Figure 1: Communication cost of anonymous broadcast protocols.

Second, can we improve performance even further by detecting and blacklisting parties that exhibit adversarial behavior? Finally, can we adopt our results to the asynchronous model of communication? We believe that this is possible for a suitably chosen upper bound on the fraction of faulty parties.

## References

- [1] M. Ajtai, J. Komlós, and E. Szemerédi. Sorting in  $c \log n$  parallel steps. *Combinatorica*, 3(1):1–19, Jan. 1983.
- [2] G. Asharov and Y. Lindell. A full proof of the BGW protocol for perfectly-secure multiparty computation. Cryptology ePrint Archive, Report 2011/136, 2011.
- [3] K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, spring joint computer conference*, AFIPS ’68 (Spring), pages 307–314, New York, NY, USA, 1968. ACM.
- [4] D. Beaver. Efficient multiparty protocols using circuit randomization. In J. Feigenbaum, editor, *Advances in Cryptology – CRYPTO ’91*, volume 576 of *Lecture Notes in Computer Science*, pages 420–432. Springer Berlin Heidelberg, 1991.
- [5] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computing. In *Proceedings of the Twentieth ACM Symposium on the Theory of Computing (STOC)*, pages 1–10, 1988.
- [6] E. Berlekamp and L. Welch. Error correction for algebraic block codes, US Patent 4,633,470, Dec. 1986.
- [7] E. Boyle, S. Goldwasser, and S. Tessaro. Communication locality in secure multi-party computation: how to run sublinear algorithms in a distributed setting. In *Proceedings of the 10<sup>th</sup> theory of cryptography conference on Theory of Cryptography*, TCC’13, pages 356–376, Berlin, Heidelberg, 2013. Springer-Verlag.
- [8] N. Braud-Santoni, R. Guerraoui, and F. Huc. Fast Byzantine agreement. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC ’13, pages 57–64, New York, NY, USA, 2013. ACM.

- [9] R. Canetti. Universally composable security: a new paradigm for cryptographic protocols. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, FOCS '01, pages 136–145, Oct 2001.
- [10] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, Feb. 1981.
- [11] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [12] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, Nov. 1998.
- [13] H. Corrigan-Gibbs and B. Ford. Dissent: accountable anonymous group messaging. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, CCS '10, pages 340–350, New York, NY, USA, 2010. ACM.
- [14] H. Corrigan-Gibbs, D. I. Wolinsky, and B. Ford. Proactively accountable anonymous messaging in verdict. In *Proceedings of the 22<sup>nd</sup> USENIX Security Symposium*, pages 147–162, Berkeley, CA, USA, 2013.
- [15] I. Damgård, Y. Ishai, M. Krøigaard, J. Nielsen, and A. Smith. Scalable multiparty computation with nearly optimal work and resilience. *Advances in Cryptology – CRYPTO '08*, pages 241–261, 2008.
- [16] V. Dani, V. King, M. Movahedi, and J. Saia. Quorums quicken queries: Efficient asynchronous secure multiparty computation. In *Distributed Computing and Networking*, volume 8314 of *Lecture Notes in Computer Science*, pages 242–256. Springer Berlin Heidelberg, 2014.
- [17] R. Dingledine, N. Mathewson, and P. Syverson. Tor: the second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, pages 21–21, Berkeley, CA, USA, 2004. USENIX Association.
- [18] S. Dolev and R. Ostrovsky. Xor-trees for efficient anonymous multicast and reception. *ACM Trans. Inf. Syst. Secur.*, 3(2):63–84, May 2000.
- [19] J. Feigenbaum and B. Ford. Seeking anonymity in an Internet panopticon. *e-Print arXiv:1312.5307*, March 2014.
- [20] S. Frank and R. Anderson. The cocaine auction protocol: On the power of anonymous broadcast. In *Proceedings of the Third International Workshop on Information Hiding*, IH 99, pages 434–447, London, UK, 2000. Springer-Verlag.
- [21] M. J. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, CCS '02, pages 193–206, New York, NY, USA, 2002. ACM.
- [22] R. Gennaro, M. Rabin, and T. Rabin. Simplified VSS and fast-track multiparty computations with applications to threshold cryptography. In *Proceedings of the 17th Annual ACM Symposium on Principles of Distributed Computing*, PODC '98, pages 101–111. ACM, 1998.
- [23] P. Golle and A. Juels. Dining cryptographers revisited. In C. Cachin and J. Camenisch, editors, *Advances in Cryptology – EUROCRYPT '04*, volume 3027 of *Lecture Notes in Computer Science*, pages 456–473. Springer Berlin Heidelberg, 2004.
- [24] M. T. Goodrich. Randomized shellsort: A simple data-oblivious sorting algorithm. *J. ACM*, 58(6):27:1–27:26, Dec. 2011.
- [25] M. T. Goodrich, M. Mitzenmacher, O. Ohrimenko, and R. Tamassia. Practical oblivious storage. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy*, CODASPY '12, pages 13–24, New York, NY, USA, 2012. ACM.
- [26] K. Hamada, R. Kikuchi, D. Ikarashi, K. Chida, and K. Takahashi. Practically efficient multi-party sorting protocols from comparison sort algorithms. In T. Kwon, M.-K. Lee, and D. Kwon, editors, *Information Security and Cryptology – ICISC 2012*, volume 7839 of *Lecture Notes in Computer Science*, pages 202–216. Springer Berlin Heidelberg, 2013.



- [27] A. Herzberg, S. Jarecki, H. Krawczyk, and M. Yung. Proactive secret sharing or: How to cope with perpetual leakage. In D. Coppersmith, editor, *Advances in Cryptology – CRYPTO '95*, volume 963 of *Lecture Notes in Computer Science*, pages 339–352. Springer Berlin Heidelberg, 1995.
- [28] J. Karlin, J. Khoury, J. Saia, and M. Zamani. Brief announcement: Scalable anonymous communication with Byzantine adversary. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 128–130, New York, NY, USA, 2013. ACM.
- [29] A. Kate, G. M. Zaverucha, and I. Goldberg. Constant-size commitments to polynomials and their applications. In *Advances in Cryptology – ASIACRYPT 2010 - 16<sup>th</sup> International Conference on the Theory and Application of Cryptology and Information Security*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, 2010.
- [30] J. Katz, C.-Y. Koo, and R. Kumaresan. Improving the round complexity of vss in point-to-point networks. In L. Aceto, I. Damgård, L. Goldberg, M. Halldorsson, A. Ingolfsson, and I. Walukiewicz, editors, *Automata, Languages and Programming*, volume 5126 of *Lecture Notes in Computer Science*, pages 499–510. Springer Berlin Heidelberg, 2008.
- [31] V. King, S. Lonargan, J. Saia, and A. Trehan. Load balanced scalable Byzantine agreement through quorum building with full information. In M. Aguilera, H. Yu, N. Vaidya, V. Srinivasan, and R. Choudhury, editors, *Distributed Computing and Networking*, volume 6522 of *Lecture Notes in Computer Science*, pages 203–214. Springer Berlin Heidelberg, 2011.
- [32] A. F. Krepinevich. Cyber warfare: A nuclear option?, 2012. Center for Strategic and Budgetary Assessments, Washington, DC, USA, 2012.
- [33] S. Laur, J. Willemson, and B. Zhang. Round-efficient oblivious database manipulation. In X. Lai, J. Zhou, and H. Li, editors, *Information Security*, volume 7001 of *Lecture Notes in Computer Science*, pages 262–277. Springer Berlin Heidelberg, 2011.
- [34] S. Le Blond, D. Choffnes, W. Zhou, P. Druschel, H. Ballani, and P. Francis. Towards efficient traffic-analysis resistant anonymity networks. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 303–314, New York, NY, USA, 2013. ACM.
- [35] T. Leighton and C. G. Plaxton. A (fairly) simple circuit that (usually) sorts. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, FOCS '90, pages 264–274, Oct 1990.
- [36] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam.  $\ell$ -diversity: Privacy beyond  $k$ -anonymity. *ACM Trans. on Knowledge Discovery from Data*, 1(1), Mar. 2007.
- [37] T. Nishide and K. Ohta. Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In T. Okamoto and X. Wang, editors, *Public Key Cryptography – PKC 2007*, volume 4450 of *Lecture Notes in Computer Science*, pages 343–360. Springer Berlin Heidelberg, 2007.
- [38] A. Pfitzmann and M. Waidner. Networks without user observability – design options. In F. Pichler, editor, *Advances in Cryptology – EUROCRYPT '85*, volume 219 of *Lecture Notes in Computer Science*, pages 245–253. Springer Berlin Heidelberg, 1986.
- [39] M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1):66–92, Nov. 1998.
- [40] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
- [41] L. von Ahn, A. Bortz, and N. J. Hopper.  $k$ -anonymous message transmission. In *Proceedings of the 10th ACM Conference on Computer and Communications Security*, CCS '03, pages 122–130, New York, NY, USA, 2003. ACM.
- [42] M. Waidner and B. Pfitzmann. The dining cryptographers in the disco: Unconditional sender and recipient untraceability with computationally secure serviceability. In J.-J. Quisquater and J. Vandewalle, editors, *Advances in Cryptology – EUROCRYPT '89*, volume 434 of *Lecture Notes in Computer Science*, page 690. Springer Berlin Heidelberg, 1990.

- [43] G. Yao and D. Feng. A new k-anonymous message transmission protocol. In C. Lim and M. Yung, editors, *Information Security Applications*, volume 3325 of *Lecture Notes in Computer Science*, pages 388–399. Springer Berlin Heidelberg, 2005.
- [44] M. Zamani, J. Saia, M. Movahedi, and J. Khoury. Towards provably-secure scalable anonymous broadcast. In *the 3rd USENIX Workshop on Free and Open Communications on the Internet*, FOCI '13, 2013.
- [45] B. Zhang. Generic constant-round oblivious sorting algorithm for MPC. In X. Boyen and X. Chen, editors, *Provable Security*, volume 6980 of *Lecture Notes in Computer Science*, pages 240–256. Springer Berlin Heidelberg, 2011.

## A Correctness and Security Proofs

In this section, we prove the correctness and secrecy of our protocols.

### A.1 Proof of Multiply

We have already showed the correctness of the **Multiply** and now prove the secrecy of the algorithm. We first define  $t$ -secrecy, a property required for the proof of secrecy.

**Definition 2.** [ $t$ -secrecy] A secret-shared value defined over a polynomial  $\phi$  is said to have  $t$ -secrecy if and only if

1. it is  $t$ -private meaning that no set of at most  $t$  parties can compute  $\phi(0)$ , and
2. it is  $t$ -resilient meaning that no set of  $t$  or less parties can prevent the other  $n - t$  remaining parties from correctly reconstructing  $\phi(0)$ .

**Lemma 5.** *The Shamir’s secret sharing scheme has the following properties.*

1. A sharing defined over a polynomial of degree  $d$  has  $d$ -secrecy if the adversary has less than  $d$  of the shares.
2. Let  $\phi_1$  and  $\phi_2$  be independent random polynomials of degree  $d_1$  and  $d_2$  that correspond to sharings with  $d_1$ -secrecy and  $d_2$ -secrecy respectively.  $\phi_3 = \phi_1 + \phi_2$  is a polynomial of degree  $d_3 = \max(d_1, d_2)$  that corresponds to a sharing with  $d_3$ -secrecy.
3. Let  $\phi_1$  and  $\phi_2$  be polynomials of degree  $d$  that correspond to two sharings both with  $d$ -secrecy.  $\phi_3 = \phi_1 \cdot \phi_2$  is a polynomial of degree  $2d$  that corresponds to a sharing with  $d$ -secrecy.

*Proof.* The first property follows from the definition. The second property is correct due to the linearity Shamir’s scheme. Without loss of generality, let  $d_1 \leq d_2$ . Intuitively, if we assume that the sharing defined by  $\phi_3$  does not have  $d_2$ -secrecy. Then, parties compute  $\phi_2 = \phi_3 - \phi_1$ . Thus, they can find  $\phi_2(0)$  (or similarly prevent others from learning  $\phi_2(0)$ ). This contradicts with the fact that  $\phi_2$  corresponds to a sharing with  $d_2$ -secrecy. For a complete proof, we refer the reader to Claim 3.4 of [2]. The third property is correct because considering an arbitrary  $P_i$  holding two shares  $a_i$  and  $b_i$  on polynomials  $\phi_1$  and  $\phi_2$  respectively,  $P_i$  learns nothing from  $a_i \cdot b_i$  other than what is revealed from  $a_i$  and  $b_i$  since  $P_i$  computes it locally. So, the resulting shared value also has  $d$ -secrecy.  $\square$

**Lemma 6.** *Let  $\langle a \rangle$  and  $\langle b \rangle$  be two secret-shared values both with  $N/3$ -secrecy and  $\langle c \rangle = \text{Multiply}(\langle a \rangle, \langle b \rangle)$ . The sharing  $\langle c \rangle$  has  $N/3$ -secrecy.*

*Proof.* In the first step of **Multiply**, algorithm **GenRand** creates two shared values  $\langle u \rangle = (u_1, \dots, u_N)$  and  $\langle v \rangle = (v_1, \dots, v_N)$  that correspond to degree  $N/6$  polynomials. Lemma 5 proves that  $\langle u \rangle$  and  $\langle v \rangle$  both have  $N/6$ -secrecy. For each party  $P_i$ , based on Lemma 5,  $w_i = u_i \cdot v_i$  defines a new sharing

that is on a polynomial of degree  $N/3$  and has  $N/6$ -secrecy. Moreover,  $\langle a \rangle$  and  $\langle b \rangle$  both are on polynomials of degree  $N/3$  and have  $N/3$ -secrecy. Thus, based on Lemma 5,  $c_i = w_i + \delta a_i + \varepsilon b_i - \varepsilon \delta$  defines a new sharing  $\langle c \rangle$  that is on a polynomial of degree  $N/3$  and has  $N/3$ -secrecy.  $\square$

**Lemma 7.** *Given two secret-shared values  $\langle a \rangle$  and  $\langle b \rangle$ , the protocol **Multiply** correctly returns a shared value  $\langle c \rangle$  such that  $c = a \cdot b$ .*

*Proof.* Follows from the proof of [4].  $\square$

## A.2 Proof of Theorem 1

In this section, we prove our main theorem (Theorem 1). We prove in the real/ideal world model as described by Canetti [9]. First, we consider the protocol in an ideal model. In this model, all parties send their input to a trusted party who computes the shuffling circuit. Then, it sends the result to all parties. Let  $A$  be the sequence of inputs and  $A'$  be the sequence of sorted inputs according to the random numbers associated with them. Recall that we have at least  $n - t$  honest parties. Based on Lemma 3 and conditioned on the event that no collision happens with high probability (Lemma 4), the elements of  $A'$  that correspond to honest parties can be any permutation of the elements of honest parties in  $A$ . In other words, the probability that the adversary can successfully map  $A'$  to  $A$  is less than  $\frac{1}{(n-t)!}$  which guarantees  $(n-t)$ -anonymity (i.e., full anonymity). Protocol 1 is the realization of the above ideal model. The real model computes the circuit in a multi-party setting. We prove this realization is correct and secure.

**Setup.** The correctness and secrecy follows from the proof of Theorem 4.

**Input Sharing.** The correctness and secrecy follows from the proof of the VSS scheme used ([30, 29]). After this step, each input quorum  $Q_i$  has a correct sharing of  $P_i$ 's input. This is the base case for our proof of circuit computation step.

**Random Generation.** The correctness and secrecy follows from the proof of the GenRand algorithm.

**Circuit Computation.** *Correctness.* We prove by induction in the real/ideal model. The invariant is that if the input shares are correct, then the output of each gate is equal to the output when the gate is computed by a trusted party in the ideal model, and the result is shared between parties of the quorum correctly. For the base case, note that the invariant is true for input gates. Induction step is based on the universal composability of **Compare** and **Multiply**. Moreover, based on the correctness of **RenewShares**, the output resharing step generates new shares without changing the output.

*Secrecy.* We prove by induction. The adversary cannot obtain any information about the inputs and outputs during the computation of each gate of the circuit. Let  $Q$ , and  $Q'$  be two quorums involved in the computation of a gate, where  $Q$  provides an input to the gate, and  $Q'$  computes the gate. Consider a party  $P$ . Let  $S$  be the set of all shares  $P$  receives during the protocol. We consider two cases. First, if  $P \notin (Q \cup Q')$ , then elements of  $S$  are independent of the shares  $Q$  sends to  $Q'$ . Moreover, elements of  $S$  are independent of the output of  $Q'$  since  $Q'$  also re-shares its output(s). Hence,  $S$  reveals nothing about the inputs and outputs of the gate.

Second, if  $P \in (Q \cup Q')$ , then the inductive invariant is that the collection of all shares held by dishonest parties in  $Q$  and  $Q'$  does not give the adversary any information about the inputs and the outputs. As the base case, it is clear that the invariant is valid for input gates. The induction step is as follows. The adversary can obtain at most  $2(N/6) = N/3$  shares of any shared value during the computation step;  $N/6$  from dishonest parties in  $Q$  and  $N/6$  from dishonest parties in  $Q'$ . By

the secrecy of the VSS scheme, at least  $N/3 + 1$  shares are required for reconstructing the secret. By the secrecy of **RenewShares** and **Multiply**, when at most  $N/3$  of the shares are revealed, the secrecy of the computation step is proved using universal computability of multi-party protocols.

**Lemma 8.** *The gate computation step of Protocol 1 obviously swaps the pairs  $(r, x)$  and  $(r', x')$  according to their first, i.e., if  $r \leq r'$ , then this step outputs  $(\langle s \rangle, \langle y \rangle) = (\langle r \rangle, \langle x \rangle)$  and  $(\langle s' \rangle, \langle y' \rangle) = (\langle r' \rangle, \langle x' \rangle)$ . Otherwise, it outputs  $(\langle s \rangle, \langle y \rangle) = (\langle r' \rangle, \langle x' \rangle)$  and  $(\langle s' \rangle, \langle y' \rangle) = (\langle r \rangle, \langle x \rangle)$ . In both cases, the adversary remains oblivious of which output is mapped to the first (second) input pair.*

*Proof.* Let  $\rho = \text{Compare}(r, r')$ . Based on the correctness and security of **Compare** (see [37]),  $\rho = 1$  if and only if  $r \leq r'$  and  $\rho = 0$ , otherwise. Also, the adversary does not learn anything about  $r$  and  $r'$ . If  $\rho = 1$ , then based on Equation 1, the linearity of the VSS scheme, and the correctness and secrecy of **Multiply** (Lemmata 6 and 7),  $s = r$ ,  $y = x$ ,  $s' = r'$ , and  $y' = x'$ . Otherwise,  $s = r'$ ,  $y = x'$ ,  $s' = r$ , and  $y' = x$ .  $\square$

**Output Propagation.** The correctness and secrecy follows from the proof of the subprotocol **Reconst**.

#### A.2.1 PROOF OF LEMMA 1 AND LEMMA 2

We first compute the cost of each step of the protocol separately and then compute the total cost. Let  $\nu_1(n)$  and  $\nu_2(n)$  be the communication and computation complexity of the VSS subprotocol respectively when it is invoked among  $n$  parties. As stated in Theorem 2 and Theorem 3 both VSS protocols used in this paper take constant rounds of communication.

- *Setup.* The communication and computation costs are equal to those costs of the quorum building algorithm of Theorem 4, which is  $\tilde{O}(1)$  for each party. This protocol takes constant rounds of communication.
- *Input Broadcast.* The input broadcast step invokes the VSS protocol  $n$  times among  $N = O(\log n)$  parties. So, this step sends  $O(n \cdot \nu_1(N))$  bits and performs  $O(n \cdot \nu_2(N))$  operations. Since the VSS scheme is constant-round, this step also takes constant rounds.
- *Random Generation.* It is easy to see that the subprotocol **GenRand** sends  $O(N \cdot \nu_1(N))$  messages, performs  $O(N \cdot \nu_2(N))$ , and has constant rounds.
- *Circuit Computation.* The sorting network of Leighton and Plaxton [35] has  $O(n \log n)$  gates. So, the communication cost of this step is equal to the communication and computation cost of running  $O(n \log n)$  instantiations of **Compare** and **RenewShares**. **Compare** requires  $O(\log q)$  invocation of **Multiply** which sends  $O(N \cdot \nu_1(N))$  messages and computes  $O(N \cdot \nu_2(N))$  operations. **RenewShares** also sends  $O(N \cdot \nu_1(N))$  messages and computes  $O(N \cdot \nu_2(N))$  operations. Hence, the circuit computation phase sends  $O(n \log n \cdot \log q \cdot N \cdot \nu_1(N))$  messages computes  $O(n \log n \cdot \log q \cdot N \cdot \nu_2(N))$ . Since the sorting network has depth  $O(\log n)$ , and **Compare** takes constant rounds, this steps takes  $O(\log n)$  rounds of communication.
- *Output Propagation.* The costs are equal to the communication and computation costs of running  $n$  invocations of **Reconst** which costs  $O(N^2)$ , plus sending the outputs to all parties, which costs  $O(n^2 N)$ . Thus, this step costs  $O(n^2 N)$ . Since **Reconst** is a constant-round protocol, this step takes constant rounds.
- *Total.* Since  $q > \frac{3}{2}kn^2 \log n$ , for a constant  $k$ ,  $q = O(n^3)$  and  $\log q = O(\log n)$ . Using eVSS, we get  $\nu_1(N) = \nu_2(N) = N^2 = O(\log^2 n)$ . Thus, Protocol 1 sends  $O(n \log^5 n)$  messages of size  $O(\kappa + \log n)$ , computes  $O(n(\kappa + \log n) \log^5 n)$  operations for shuffling  $n$  messages (excluding

the output step). This proves Lemma 2. For Lemma 1, since  $\nu_1(N) = \nu_2(N) = O(\text{poly}(N))$ , Protocol 1 sends  $\tilde{O}(n)$  bits and computes  $\tilde{O}(n)$  operations for shuffling  $n$  messages. This proves Lemma 1. In both cases, the output propagation step costs  $O(n^2 \log n)$  field elements. Finally, in both cases, the protocol requires  $O(\log n)$  rounds of communication. This finishes the proof of Theorem 1.

## B Remaining Algorithms

Beaver [4] describes a simple technique for generating uniform random secrets by adding shared values chosen uniformly at random by each party. Such a random shared value is used in several parts of our protocol. The following subprotocol implements this technique.

---

### Protocol 3 GenRand

---

Usage. Parties jointly generate a shared value  $\langle r \rangle = (r_1, \dots, r_n)$  where  $r$  is chosen uniformly at random from  $\mathbb{Z}_p$ .

GenRand():

For all  $i \in [n]$ ,

1. Party  $P_i$  chooses  $\rho_i \in \mathbb{Z}_p$  uniformly at random and secret shares it among all parties.
  2. Let  $\rho_{1i}, \dots, \rho_{ni}$  be the shares  $P_i$  receives from the previous step.  $P_i$  computes  $r_i = \sum_{j=1}^n \rho_{ji}$ .
- 

In the malicious setting, it is possible that dishonest parties send spurious shares during secret reconstruction phase. In eVSS [29] (used in the cryptographic version of our protocol), this is solved by asking all parties to broadcast a proof (called *witness*) during reconstruction to verify broadcast shares. In our protocol, reconstruction is postponed to after circuit computation. Since the witnesses are generated in the sharing phase at the beginning of the computation, and the witnesses do not have necessary homomorphic properties, we cannot use them in our reconstruction phase. Instead, we correct corruptions using a BCH decoding algorithm (e.g., the algorithm of Berlekamp and Welch [6]) as in normal secret reconstruction [4]. This technique is also used in the VSS of Katz et al. [30] and is implemented in the following subprotocol.

---

### Protocol 4 Reconst

---

Usage. Initially, all parties jointly hold a shared value  $\langle a \rangle = (a_1, \dots, a_n)$ . Using this algorithm, parties jointly reconstruct the secret, i.e., all parties learn the value  $a$ .

Reconst( $\langle a \rangle$ ):

For all  $i \in [n]$ ,

1. Party  $P_i$  sends its share  $a_i$  to all parties via one-to-one communication.
2. Let  $\alpha_1, \dots, \alpha_n$  be the messages  $P_i$  receives from the previous step.  $P_i$  computes a degree  $d = n/3$  polynomial  $\phi(x)$  using the Lagrange interpolation polynomial,

$$\phi(x) = \sum_{i=1}^{d+1} \alpha_i \prod_{j=1, j \neq i}^{d+1} (x - j)(i - j)^{-1}$$

3. For all  $j \in [n]$ , if there exists at least one  $\alpha_j$  such that  $\phi(j) \neq \alpha_j$ , then  $P_i$  runs the decoding algorithm of Berlekamp and Welch [6] to recover the correct polynomial  $\phi'(x)$  of degree  $d$ . For all  $j \in [k]$ , if  $\phi'(j) \neq \alpha_j$ , then  $P_i$  concludes that  $P_j$  is dishonest and must be disqualified.
-